



Microsoft

70-765 Exam

Microsoft Provisioning SQL Databases Exam

Thank you for Downloading 70-765 exam PDF Demo

You can Buy Latest 70-765 Full Version Download

<https://www.certkillers.net/Exam/70-765>

<https://www.certkillers.net>

Version: 21.1

Question: 1

You have a Microsoft SQL Server 2014 named SRV2014 that has a single tempdb database file. The tempdb database file is eight gigabytes (GB) in size.

You install a SQL Server 2016 instance named SQL Server 2016 by using default settings. The new instance has eight logical processor cores.

You plan to migrate the databases from SRV2014 to SRV2016.

You need to configure the tempdb database on SRV2016. The solution must minimize the number of future tempdb autogrowth events.

What should you do?

A. Increase the size of the tempdb datafile to 8 GB. In the tempdb database, set the value of the MAXDOP property to 8.

B. Increase the size of the tempdb data files to 1 GB.

C. Add seven additional tempdb data files. In the tempdb database, set the value of the MAXDOP property to 8.

D. Set the value for the autogrowth setting for the tempdb data file to 128 megabytes (MB). Add seven additional tempdb data files and set the autogrowth value to 128 MB.

Answer: B

Explanation:

In an effort to simplify the tempdb configuration experience, SQL Server 2016 setup has been extended to configure various properties for tempdb for multi-processor environments.

1. A new tab dedicated to tempdb has been added to the Database Engine Configuration step of setup workflow.

2. Configuration options:

Data Files

* Number of files – this will default to the lower value of 8 or number of logical cores as detected by setup.

* Initial size – is specified in MB and applies to each tempdb data file. This makes it easier to configure all files of same size. Total initial size is the cumulative tempdb data file size (Number of files * Initial Size) that will be created.

* Autogrowth – is specified in MB (fixed growth is preferred as opposed to a non-linear percentage based growth) and applies to each file. The default value of 64 MB was chosen to cover one PFS interval.

Figure:

SQL Server 2016 CTP3.0 Setup

Database Engine Configuration

Specify Database Engine authentication security mode, administrators, data directories and TempDB settings.

Product Key
License Terms
Global Rules
Product Updates
Install Setup Files
Install Rules
Installation Type
Setup Role
Feature Selection
Feature Rules
Instance Configuration
Server Configuration
Database Engine Configuration
Feature Configuration Rules
Ready to Install
Installation Progress
Complete

Server Configuration | Data Directories | **TempDB** | FILESTREAM

TempDB data files: tempdb.mdf, tempdb_mssql_#.ndf

Number of files: 4

Initial size (MB): 8 Total initial size (MB): 32

Autogrowth (MB): 64 Total autogrowth (MB): 256

Data directories: C:\Program Files\Microsoft SQL Server\MSSQL13.SQL16CTP3\M... Add... Remove

TempDB log file: templog.ldf

Initial size (MB): 8

Autogrowth (MB): 64

Log directory: C:\Program Files\Microsoft SQL Server\MSSQL13.SQL16CTP3\M... ..

< Back Next > Cancel Help

References:

<https://blogs.msdn.microsoft.com/psssql/2016/03/17/sql-2016-it-just-runs-faster-automatic-tempdb-configuration/>

Question: 2

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You have deployed several GS-series virtual machines (VMs) in Microsoft Azure. You plan to deploy Microsoft SQL Server in a development environment. Each VM has a dedicated disk for backups. You need to backup a database to the local disk on a VM. The backup must be replicated to another region.

Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 disk storage
- C. Premium P30 disk storage
- D. Standard locally redundant disk storage
- E. Standard geo-redundant disk storage
- F. Standard zone redundant blob storage

- G. Standard locally redundant blob storage
- H. Standard geo-redundant blob storage

Answer: E

Explanation:

Note: SQL Database automatically creates a database backups and uses Azure read-access geo-redundant storage (RA-GRS) to provide geo-redundancy. These backups are created automatically and at no additional charge. You don't need to do anything to make them happen. Database backups are an essential part of any business continuity and disaster recovery strategy because they protect your data from accidental corruption or deletion.

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-automated-backups>

Question: 3

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You have a virtual machine (VM) in Microsoft Azure, which has a 2 terabyte (TB) database. Microsoft SQL Server backups are performed by using Backup to URL.

You need to provision the storage account for the backups while minimizing costs.

Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 disk storage
- C. Premium P30 disk storage
- D. Standard locally redundant disk storage
- E. Standard geo-redundant disk storage
- F. Standard zone redundant blob storage
- G. Standard locally redundant blob storage
- H. Standard geo-redundant blob storage

Answer: G

Explanation:

A URL specifies a Uniform Resource Identifier (URI) to a unique backup file. The URL is used to provide the location and name of the SQL Server backup file. The URL must point to an actual blob, not just a container. If the blob does not exist, it is created. If an existing blob is specified, BACKUP fails, unless the "WITH FORMAT" option is specified to overwrite the existing backup file in the blob. LOCALLY REDUNDANT STORAGE (LRS) makes multiple synchronous copies of your data within a single datacenter.

Question: 4

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You have deployed a GS-series virtual machine (VM) in Microsoft Azure. You plan to deploy Microsoft SQL Server.

You need to deploy a 30 megabyte (MB) database that requires 100 IOPS to be guaranteed while minimizing costs.

Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 disk storage
- C. Premium P30 disk storage
- D. Standard locally redundant disk storage
- E. Standard geo-redundant disk storage
- F. Standard zone redundant blob storage
- G. Standard locally redundant blob storage
- H. Standard geo-redundant blob storage

Answer: A

Explanation:

Premium Storage Disks Limits

When you provision a disk against a Premium Storage account, how much input/output operations per second (IOPS) and throughput (bandwidth) it can get depends on the size of the disk. Currently, there are three types of Premium Storage disks: P10, P20, and P30. Each one has specific limits for IOPS and throughput as specified in the following table:

Premium Storage Disk Type	P10	P20	P30
Disk Size	128 GiB	512 GiB	1024 GiB (1 TB)
IOPS per disk	500	2300	5000
Throughput per disk	100 MB per second	150 MB per second	200 MB per second

References:

<https://docs.microsoft.com/en-us/azure/storage/storage-premium-storage>

Question: 5

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent

of the other questions in this series. Information and details provided in a question apply only to that question.

You have deployed several GS-series virtual machines (VMs) in Microsoft Azure. You plan to deploy Microsoft SQL Server in a development environment.

You need to provide storage to the environment that minimizes costs.

Which storage option should you use?

- A. Premium P10 disk storage
- B. Premium P20 disk storage
- C. Premium P30 disk storage
- D. Standard locally redundant disk storage
- E. Standard geo-redundant disk storage
- F. Standard zone redundant blob storage
- G. Standard locally redundant blob storage
- H. Standard geo-redundant blob storage

Answer: D

Question: 6

HOTSPOT

You use Resource Manager to deploy a new Microsoft SQL Server instance in a Microsoft Azure virtual machine (VM) that uses Premium storage. The combined initial size of the SQL Server user database files is expected to be over 200 gigabytes (GB). You must maximize performance for the database files and the log file.

You add the following additional drive volumes to the VM:

Drive volume	Storage	Host caching
E:	Premium storage	ReadOnly
F:	Premium storage	None

You have the following requirements:

You need to deploy the SQL instance.

In the table below, identify the drive where you must store each SQL Server file type.

NOTE: Make only one selection in each column. Each correct selection is worth one point.

Answer area

Drive	Data files	Log files
C:	<input type="radio"/>	<input type="radio"/>
D:	<input type="radio"/>	<input type="radio"/>
E:	<input type="radio"/>	<input type="radio"/>
F:	<input type="radio"/>	<input type="radio"/>

Answer:

Answer area

Drive	Data files	Log files
C:	<input type="radio"/>	<input type="radio"/>
D:	<input type="radio"/>	<input type="radio"/>
E:	<input checked="" type="radio"/>	<input type="radio"/>
F:	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

Enable read caching on the disk(s) hosting the data files and TempDB.

Do not enable caching on disk(s) hosting the log file. Host caching is not used for log files.

Question: 7

DRAG DROP

You are building a new Always On Availability Group in Microsoft Azure. The corporate domain controllers (DCs) are attached to a virtual network named ProductionNetwork. The DCs are part of an availability set named ProductionServers1.

You create the first node of the availability group and add it to an availability set named ProductionServers2. The availability group node is a virtual machine (VM) that runs Microsoft SQL

Server. You attach the node to ProductionNetwork.

The servers in the availability group must be directly accessible only by other company VMs in Azure.

You need to configure the second SQL Server VM for the availability group.

How should you configure the VM? To answer, drag the appropriate configuration settings to the correct target locations. Each configuration setting may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Configuration settings

None/Not Assigned

ProductionServers1

ProductionNetwork

ProductionServers2

Create a new Object

VM settings page

Settings

Storage

Disk type

StandardPremium (SSD)

* Storage account (new) sqlstorage3

Network

* Virtual network setting

* Subnet ProductionServers (10.1.0.0/24)

* Public IP address setting

* Network security group (new) SQLServers

Extensions

Extensions No extensions

Monitoring

Diagnostics DisabledEnabled

Availability

* Availability set setting

OK

Answer:

VM settings page

Settings — □ X

Storage

Disk type !

Standard Premium (SSD)

* Storage account ! >

(new) sqlstorage3

Network

* Virtual network ! >

ProductionNetwork >

* Subnet ! >

ProductionServers (10.1.0.0/24)

* Public IP address ! >

None/Not Assigned >

* Network security group ! >

(new) SQLServers

Extensions

Extensions ! >

No extensions

Monitoring

Diagnostics !

Disabled Enabled

Availability

* Availability set ! >

ProductionServers2 >

Explanation:

Box 1: ProductionNetwork

The virtual network is named ProductionNetwork.

Box 2: None /Not Assigned

As the servers in the availability group must be directly accessible only by other company VMs in Azure, there should be no Public IP address.

Box 3: ProductionServer2

You create the first node of the availability group and add it to an availability set named ProductionServers2. The availability group node is a virtual machine (VM) that runs Microsoft SQL Server.

Question: 8

HOTSPOT

You plan to migrate a Microsoft SQL Server workload from an on-premises server to a Microsoft Azure virtual machine (VM). The current server contains 4 cores with an average CPU workload of 6 percent and a peak workload of 10 percent when using 2.4Ghz processors.

You gather the following metrics:

	Minimum IOPS	Average IOPS	Maximum IOPS
Data Drive	100	938	7253
Transaction Log Drive	12	145	350
TempDB Drive	300	900	1900

You need to design a SQL Server VM to support the migration while minimizing costs.

For each setting, which value should you use? To answer, select the appropriate storage option from each list in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

VM setting	Value ▼
Data drive	Local storage Premium storage Standard storage
Transaction log drive	Local storage Premium storage Standard storage
TempDB drive	Local storage Premium storage Standard storage
VM size	A3 D3 DS3

Answer:

Answer Area

VM setting	Value ▼
Data drive	Local storage Premium storage Standard storage
Transaction log drive	Local storage Premium storage Standard storage
TempDB drive	Local storage Premium storage Standard storage
VM size	A3 D3 DS3

Explanation:

Data drive: Premium Storage

Transaction log drive: Standard Storage

TempDB drive: Premium Storage

Note: A standard disk is expected to handle 500 IOPS or 60MB/s.

A P10 Premium disk is expected to handle 500 IOPS.

A P20 Premium disk is expected to handle 2300 IOPS.

A P30 Premium disk is expected to handle 5000 IOPS.

VM size: A3

Max data disk throughput is 8x500 IOPS

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/virtual-machines-windows-sizes>

Question: 9

You plan to migrate a database To Microsoft Azure SQL Database. The database requires 500 gigabytes (GB) of storage.

The database must support 50 concurrent logins. You must minimize the cost associated with hosting the database.

You need to create the database.

Which pricing tier should you use?

- A. Standard S3 pricing tier
- B. Premium P2tier
- C. Standard S2 pricing tier
- D. Premium P1 tier

Answer: D

Explanation:

For a database size of 500 GB the Premium tier is required.

Both P1 and P2 are adequate. P1 is preferred as it is cheaper.

Note:

Premium service tier

Service tier	P1	P2	P4	P6	P11	P15
Max DTUs	125	250	500	1000	1750	4000
Max database size*	500 GB	500 GB	500 GB	500 GB	1 TB	1 TB
Max in-memory OLTP storage	1 GB	2 GB	4 GB	8 GB	14 GB	32 GB
Max concurrent workers	200	400	800	1600	2400	6400
Max concurrent logins	200	400	800	1600	2400	6400
Max concurrent sessions	30000	30000	30000	30000	30000	30000

Question: 10

HOTSPOT

You need to ensure that a user named Admin2 can manage logins.

How should you complete the Transact-SQL statements? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

▼	Admin2 WITH password = 'Pa\$\$w0rd';	
CREATE USER		
ALTER SERVER ROLE		
CREATE LOGIN		

▼	Admin2User FROM	▼	Admin2
CREATE USER		WINDOWS	
ALTER SERVER ROLE		EXTERNAL PROVIDER	
CREATE LOGIN		LOGIN	

ALTER ROLE '	▼
	loginmanager
	dbmanager
	bd_ddladmin

Answer:

The screenshot shows the SQL Azure portal interface for creating a user and login. It is divided into three main sections:

- Top Section:** A dropdown menu is set to 'Admin2' with the password 'Pa\$\$w0rd'. The 'CREATE LOGIN' option is highlighted with a red box.
- Middle Section:** A dropdown menu is set to 'Admin2User' with the text 'FROM'. To the right, another dropdown menu is set to 'Admin2' with the 'LOGIN' option highlighted with a red box.
- Bottom Section:** Labeled 'ALTER ROLE', it shows a dropdown menu with 'loginmanager' highlighted with a red box. Other options visible are 'dbmanager' and 'bd_ddladmin'.

Explanation:

Step 1: CREATE LOGIN

First you need to create a login for SQL Azure, it's syntax is as follows:

CREATE LOGIN username WITH password='password';

Step 2, CREATE USER

Step 3: LOGIN

Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database. Here is some sample Transact-SQL that creates a user:

CREATE USER readonlyuser FROM LOGIN readonlylogin;

Step 4: loginmanager

Members of the loginmanager role can create new logins in the master database.

References:

<https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/>

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-manage-logins>

Thank You for trying 70-765 PDF Demo

To Buy Latest 70-765 Full Version Download visit link below

<https://www.certkillers.net/Exam/70-765>

Start Your 70-765 Preparation

[Limited Time Offer] Use Coupon “CKNET” for Further discount on your purchase. Test your 70-765 preparation with actual exam questions.

<https://www.certkillers.net>